

Claims

The invention claimed is:

1. A method of managing data movement, comprising:

establishing a processing environment in a cluster of nodes having common access to data residing in one or more data storage units;

initiating a data management application (DM) in said environment;

assigning a node of said cluster as a coordinating node for managing data movement;

receiving an event by the coordinating node requesting movement of data;

posting a worker thread to one or more of the nodes to perform data movement in response to the event.

2. The method of claim 1, wherein said worker threads are posted to one or more nodes other than said coordinating node to perform data movement tasks.

3. The method of claim 1, wherein said coordinating node is a session node.

4. The method of claim 1, further comprising providing data management access rights to the one or more nodes to which said worker threads are posted, and permitting only the one or more nodes having said data management access rights to execute said worker threads.

5. The method of claim 1, further comprising establishing a process session in said cluster and assigning a session identifier for that session.

6. The method of claim 5, further comprising providing said session identifier to said one or more nodes to which said worker threads are posted, and permitting only the one or more nodes having said session identifier to execute said worker thread.

7. The method of claim 5, wherein said DM application establishes said session and assigns said session identifier.

8. The method of claim 5, wherein a plurality of sessions are established in said cluster concurrently and each session is assigned a unique session identifier.

9. The method of claim 1, wherein said DM application utilizes one or more parallel file systems for management of data.

10. The method of claim 9, wherein each parallel file system further comprises one or more physical file systems.

11. The method of claim 10, wherein said worker threads include calls for performing at least one of punching holes in files, moving data into files and moving data out of files.

12. The method of claim 9, wherein said DM application is initiated using a data management application programming interface (DMAPI).

13. The method of claim 1, wherein said DM application is initiated using a data management application programming interface (DMAPI).

14. The method of claim 1, wherein said processing environment includes a storage area network (SAN) including said one or more data storage units.

15. The method of claim 12, wherein said processing environment includes a storage area network (SAN) including said one or more data storage units.

16. The method of claim 14, wherein said worker threads perform data movement within a hierarchical storage management (HSM) system.

17. The method of claim 1, further comprising reassigning a worker thread to another node upon failure of the node to which the worker thread is dispatched.

18. The method of claim 1, further comprising assigning another coordinating node upon failure of the coordinating node.

19. A machine readable medium having a set of instructions recorded thereon for performing a method of managing data movement, said method including:

establishing a processing environment in a cluster of nodes having common access to data residing in one or more data storage units;

initiating a data management application (DM) in said environment;

assigning a node of said cluster as a coordinating node for managing data movement;

receiving an event by the coordinating node requesting movement of data;

posting a worker thread to one or more of the nodes to perform data movement in response to the event.

20. A system for managing data movement comprising;

a computing environment having a cluster of nodes having common access to data residing in one or more data storage units;

a data management application (DM) operable to manage data movement by assigning any node in said cluster as a coordinating node to manage data movement events and dispatching worker threads to one or more nodes to perform data movement tasks in response to the data movement events.